arXiv:1604.04048v1 [cs.CV] 14 Apr 2016

Deep Feature Based Contextual Model for Object Detection

Wenqing Chu and Deng Cai

State Key Lab of CAD & CG, Zhejiang University

Abstract. Object detection is one of the most active areas in computer vision, which has made significant improvement in recent years. Current state-of-the-art object detection methods mostly adhere to the framework of regions with convolutional neural network (R-CNN) and only use local appearance features inside object bounding boxes. Since these approaches ignore the contextual information around the object proposals, the outcome of these detectors may generate a semantically incoherent interpretation of the input image. In this paper, we propose an ensemble object detection system which incorporates the local appearance, the contextual information in term of relationships among objects and the global scene based contextual feature generated by a convolutional neural network. The system is formulated as a fully connected conditional random field (CRF) defined on object proposals and the contextual constraints among object proposals are modeled as edges naturally. Furthermore, a fast mean field approximation method is utilized to inference in this CRF model efficiently. The experimental results demonstrate that our approach achieves a higher mean average precision (mAP) on PASCAL VOC 2007 datasets compared to the baseline algorithm Faster R-CNN.

Keywords: Object Detection, Context Information, Conditional Random Field

1 Introduction

Object detection is one of the fundamental problems in computer vision. It plays an important role in many real-world applications such as image retrieval, advanced driver assistance system and video surveillance. This problem is very difficult because the object appearances vary dramatically from changes in different illuminations, view points, nonrigid deformations, poses, and the presence of occlusions. For instance, there is a large amount of partial occlusions between pedestrians standing next to each other in a crowd street and birds come in various poses and colors.

In the past few years, remarkable progress has been made to boost the performance of object detection. A common pipeline to address this problem consists of two main steps: (1) object proposal generation, (2) class-specific scoring and bounding box regression. There is a significant body of methods for generating

object proposals such as [1,2,3,4,5] or just a sliding window fashion [6]. Then some specific feature of the object bounding box is extracted and some classifier is applied for efficient and accurate object detection, in which the representative work includes AdaBoost algorithm [7], DPM model [6] and deep CNN models [8]. However, most state-of-the-art detectors like Faster RCNN [9] only consider the proposals individually without taking the contextual information into account.

In the real world, there exists a semantic coherent relationship between the objects both in terms of relative spatial location and co-occurrence probability [10,11]. In some situations, contextual information among objects in the input image can provide a more valuable cue for the detection of an object than the information near an objects region of interest. In addition, the global context based on scene understanding also helps the detector better rule out the false alarm.

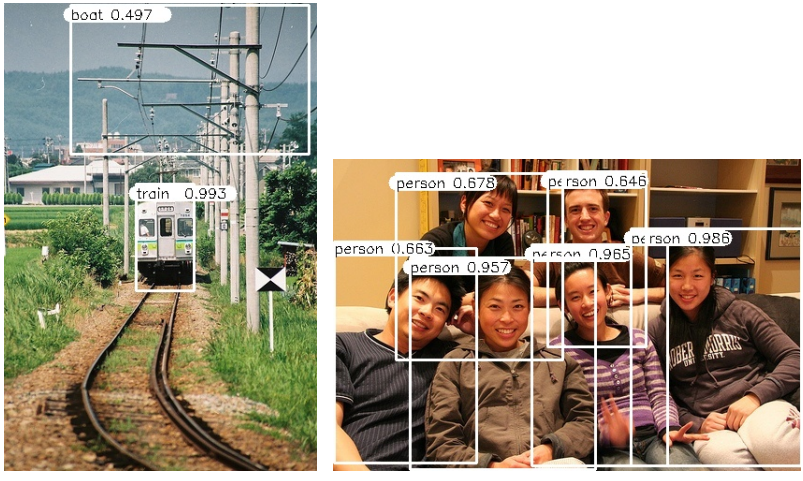

(a) Boat and Train (b) Partial Occlusions between people

Fig. 1. Object-level contextual information

In Figure 1(a), Faster R-CNN recognizes the two object proposals individually. However, boats and trains stand little chance of co-occurence in the input image , which means the probability for the boat or the train should be decreased. In Figure 1(b), since the person is occluded by the sofa and Faster R-CNN classify the object proposal into person by the probability of 0.646. However, if we can use the contextual information of objects around the bounding box to help inference, we can raise our confidence that the category of this object candidate behand the sofa is a person. In addition, image-level contextual information can also support object detection. If we can recognize the scene in Figure 2(a) as a lake or sea, then we can make a better judgement that the object proposals are

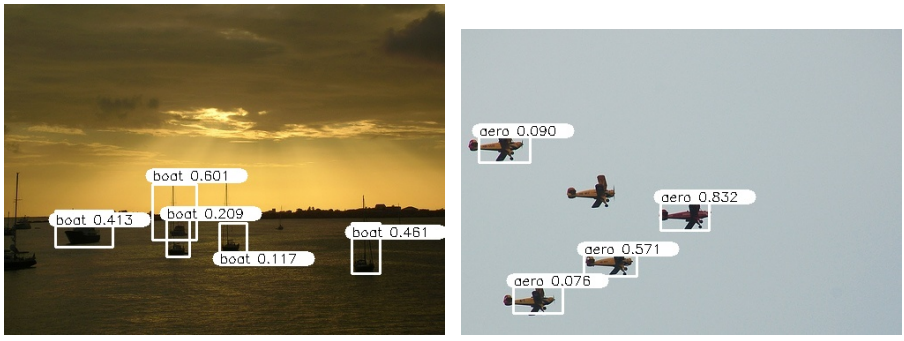

(a) Boats in Lake (b) Aeroplanes in air

Fig. 2. Image-level contextual information

likely to be tiny boats. Similarily, apperance in Figure 2(b) looks like sky which enhances the probability of the presence of aeroplanes. In a word, the context information can be a strong clue for recognizing the object candidates which are ambiguous because of low resolution and variation in pose and illumination.

In the history of computer vision, a number of approaches [12,13,14,15,16,17,18,19,20] have exploited contextual information in order to boost the preformance of object detection. Nevertheless, most of these methods leverage hand-crafted features such as Gist [12] or HOG [21] to extract features from the input image. Recently, convolutional neural network (CNN) has achieved great success in computer vision tasks such as image classification [22], which inspired us to employ CNN to devise a novel contextual model.

In this paper, we propose a novel context model which is based on the prominent object detection method Faster R-CNN [9] in recent years. To leverage the context information around each candidate, we first focus on the local contextual classes that are present near the object proposal. As in [16], inter-object constraints vary greatly from changes in different categories and locations, which can be learned from the statistical summary of the datasets. In addition, we also generate the global scene descriptor using a CNN model which is trained for the secne understanding task. Then the global scene information of the input image is utilized as the input of a logistic regression method to predict the probability how much some category is likely to occur in the image. In the following step, we apply Faster R-CNN to each input image and obtain a pool of object proposals with the corresponding scores and locations for each category. After that, we take these object proposals as nodes in graph and formulate the graph as a fully-connected conditional random field (CRF) according to the context information. Specifically, the unary potentials are determined by the scores from Faster R-CNN. And the pairwise postentials are decided by the layouts and categories among the object proposals. To efficiently inference in this CRF, we utilize a

fast mean field inference algorithm of [23,24] to yield the candidate labels and the corresponding confidence simultaneously.

We have extensively evaluated our method on the PASCAL VOC2007 dataset [25]. The experimental results show that our method can acheieve an improvement of 0.87% in mAP and for the category bottle the AP can be raised up to 3.4%. Very few categories accuracies are worsened by context information.

The rest of this paper is arranged as follows. First, we briefly review a few of recent work on object detection methods in Section 2. Then in section 3 we describe the framework of our context model for object detection and the inference algorithm in detail. After that, we evaluate the performance of our method on the challenging databases PASCAL 2007 in section 4. Finally, in section 5, we present our conclusions and discuss future work.

2 Related Work

In this section, we briefly review the recent work on object detection. Object detection has been active research areas in recent years, which has lead to a large amount of methods to address the problems in it.

In the literature of object detection, the part-based model is one of the most powerful approaches in which deformable part-based model (DPM) [6] is an excellent example. This method utilize mixtures of multiscale deformable part models to represent highly variable objects. Each part captures local appearance properties of an object while the deformable configuration is characterized by spring-like connections between certain pairs of parts. It detect objects in an image by a sliding window approach and take the histogram of oriented gradients (HOG) features [21] as input.

Recently, deep convolutional neural networks (CNN) have emerged as a powerful machine learning model on a number of image recognition benchmarks, including the most noticeably work by [22]. That aroused a significant body of methods [26,27,28,8,29,30,9,31,32,33,34,35] addressing the problem with CNN. Among these approaches, the regions-with-convolutional-neural-network (R-CNN) framework [8] achieved excellent detection performance and becomes a commonly employed paradigm for object detection. Its essential steps include proposal generation with selective search [2], CNN feature extraction, object box classification and regression based on the CNN feature. However, R-CNN brings excessive computation cost because it extracts CNN feature repeatedly for thousands of object proposals. Spatial pyramid pooling networks (SPPnets) [29] were proposed to accelerate the process of feature extraction in R-CNN by sharing the forward pass computation.The SPPnet approach computes a convolutional feature map for the entire input image once and then generates a fixed-length feature vector from the shared feature map for each object proposal. Fast Region-based Convolutional Network method (Fast R-CNN) [30] utilizes a multi-task loss,which leads to an end-to-end framework that the training is a single-stage and no disk storage is required for feature caching. The drawback of Fast R-CNN is that this method still use bottom-up proposal generation which is the

bottleneck of efficiency. In [9], the authors proposed a Region Proposal Network (RPN) method that shares full-image convolutional features with the detection network, thus enabling nearly cost-free region proposals.These techniques, however, still mostly perform detection based on local appearance features in the input image.

In the other hand, semantic context also plays a very important role of boosting the performance of object detection [12,14,16,18,19,20]. The statistics of low-level features across the entire scene were used to predict the presence or absence of objects and their locations [12]. In [14], the authors demonstrated that contextual relations between objects' labels can help reduce ambiguity in objects' visual appearance. Specifically, they utilized image segmentation as a pre-processing step to generate object proposals. Then a conditional random field (CRF) formulation was exploited as post-processing to infer the optimal label configuration of this CRF model, which jointly labels all the object candidates. [16] extend this approach to combine two types of context co-occurrence and relative location with local appearance based features. [19] introduced a unified model for multi-class object recognition that learns statistics that capture the spatial arrangements of various object classes in real images.

Besides, some work [36,37,38] which focus on detecting some specialized object class were proposed to use context information to support detection. However, these models mostly work on context information represented by traditional visual features such as HOG or GIST. Thus, we are motivated to move on to more powerful features provided by CNN model.

3 A Fully-Connected CRF for Object Detection

In this section, we address the general object detection problem with an ensemble system, in which we combine the local appearence, the contextual relationships among object candidates and the global scene context information. To process each input image, our approach includes three main stages. At first, we generate a pool of object proposals obtained with Faster R-CNN [9]. Then we use a conditional random field (CRF) framework to model the object detection problem with contextual information. Finally, we employ an efficient mean field approximation method to inference and finaly maximize object label agreement.

In Section 3.1, we introduce the process of object proposal generation building on Faster R-CNN [9]. In Section 3.2, We give the formulation of the CRF model for the object detection problem. We will describe unary potentials, pairwise potentials between object candidates and the global potentials determined by the context feature describing the entire image. After all, Section 3.3 will give the inference algorithm.

3.1 Object Proposals

Our approach generates object proposals following Faster R-CNN [9], which is one of the state-of-the-art object detection methods. In contrast to R-CNN [8],

Faster R-CNN proposed a Region Proposal Network (RPN) instead of other bottom-up approaches to outputs a set of object bounding boxes. Since RPN slides a small network over the conv feature map output by the last conv layer, which makes it can share forward pass computation with a Fast R-CNN object detection network [30] and that leads to great advantages on efficiency.

Faster R-CNN method is a object detection system which can depend on different CNN architecture. In our experiments, we investigate the Zeiler and Fergus model [39] (ZF), which has 5 shareable conv layers and the Simonyan and Zisserman model [40] (VGG), which has 13 shareable conv layers. To verify our method is insensitive to different object proposal methods, we conduct experiments on both CNN models. To train the network, we optimize parameters with the popular stochastic gradient descent (SGD) [41] with momentum. The Faster R-CNN model is pre-trained on the ImageNet dataset [42] and fintuned on the PASCAL VOC 2007 dataset [25]. More details on the training procedure can be found in [9].

3.2 CRF Formulation

In this stage, we take each object proposal generated by the first step in the pipeline as a node. Suppose there are N proposals in a single input image and K categories, we can consider a random field $\mathbf{X}$ defined over a set of variables $\{x_1, \ldots, x_N\}$. The domain of each variable is a set of categories $L = \{0, 1, 2, \ldots, K\}$ in which 0 represents background and x_j is the label assigned to proposal j. Faster R-CNN generates a scores matrix $\mathbf{S} \in \mathbb{R}^{N \times K}$ which $\mathbf{S}_{j,k}$ represents the probability of proposal j belong to category. These scores are used as initial scores of each node. Our method will adjust the scores of each node in a fully connected graph according to the scene information and the contextual relationship among the nodes.

Unary Potentials In our conditional random field (CRF) model, the unary potential $\phi_u(x_i)$ measures the probability that the object proposal i belongs to the category x_i according to the local appearance. Therefore, we use a rescaled score from Faster R-CNN as unary potential. This lets us write our unary term as

$$\phi_u(x_i) = -log(\mathbf{S}_{i,x_i}) \tag{1}$$

where $\mathbf{S}_{i,x_i}$ is the confidence that proposal i belong to category x_i.

Pairwise Potentials Here we describe our pairwise potentials that purpose to capture the contextual information between multiple object candidates. Our pairwise model takes both the semeantic and spatial relatioinships into account like [16,19]. In our method, we define a function $\phi_p()$ to estimate the pairwise potential of a proposal i based on a known object j in the same image. The input of $\phi_p()$ is the category information of i and j and their relative location. To leverage this spatial relevance better, we consider 11 different layouts for

two object bounding boxes. If two candidates dont't have intersection, then the spatial relationships of them can be classified into far, up, down, left and right. Otherwise, they can be defined by inside, outside, up, down, left and right. For example, some pottedplants are above another in Figure 3(a) and don't intersect between each other. InFigure 3(b), however, the human is on the horse and their bounding boxes have intersection. We define these situations as two different spatial relationships.

(a) Pottedplants (b) Horse and Human

Fig. 3. different spatial relationship

Following [16,19], this function can be defined according to the likelihood $P(x_i, x_j, r)$ which is learned from statistic summary of the training dataset.

$$\phi_p(x_i, x_j, r) = -log(P(x_i, x_j, r)) \tag{2}$$

where $P(x_i, x_j, r)$, with r = 1,...,11, measures the likelyhood that an object with label x_i appears with an object label x_j for a given relationship r.

Global Potentials In addition to object-level contextual information, we also introduce image-level signal to reason about presence or absence of objects in the image. The key point is to find global image features which can represent various secne well.

Scene categorization is also a challenge task in computer vision. Before, most work focus on shallower hand-crafted features empirically and the databases that they used are lack of abundance and variety. Recently, the Places2 dataset [43] is provided which contains more than 10 million images comprising 400+ unique scene categories. Moveover, the dataset features 5000 to 30,000 training images per class, consistent with real-world frequencies of occurrence. With this large dataset, we can apply the powerful CNN model to feature extraction representing the scene context information. Specifically, the CNN model takes the whole image as input and outputs a score for each category. Following [44], we train the VGG network [40] on the dataset using Caffe deep learning toolbox. And the last layer

is used to represent the scene context. With the generic deep scene features for visual recognition, we use a logistic regression model to fit it and output $p(x_i|f)$ which measures the probability of existence of the category x_i in the input image.

$$\phi_g(x_i) = -log(p(x_i|f)) \tag{3}$$

where $(x_i|f)$ measures the likelyhood that an object with label x_i appears in the input image.

Then we formulate the fully-connected CRF model as below:

$$E(\mathbf{X}) = \sum_{i=1}^{N} \phi_u(x_i) + \omega_p \sum_{i<j} \phi_p(x_i, x_j, r) + \omega_g \sum_{i=1}^{N} \phi_g(x_i) \tag{4}$$

where ω_p and $\omega_g$are the weights of the pairwise potentials and the global potentials.

3.3 Inference Algorithm

To tackle this fully-connected CRF model, we use the mean field approximation method to minimize the objective function. Following [23], we adopt a fast mean field approximation algorithm to compute the marginals. Given the current mean field estimates $\{Q_i\}$ of the marginals, the update equation can be written as

$$Q_i(x_i) \propto exp(-\phi_u(x_i) - \omega_g \phi_g(x_i) - \omega_p \sum_{x_j \neq x_i} \sum_{j \neq i} Q_j(x_j)\phi_p(x_j, x_i)) \tag{5}$$

After convergence, we obtain an (approximate) posterior distribution of object labels for each node. To obtain the final results, we can employ the mean field approximate marginal probability $Q_i(x_i)$ as a detection score. Since the number of object proposals is mostly around 300, the time cost is almost free. The whole procedure for inference has been presented in Algorithm 1..

4 Experiments

In this section, we conduct a series of experiments to evaluate the performance of our approach and compare it against the state-of-the-art object detection baseline. We evaluate our method on object detection benchmark datasets PASCAL VOC 2007 [25]. There are 20 different categories of objects and every dataset is divided into train, val and test subsets. In this dataset, object appearances vary greatly from changes in different illuminations, poses, locations, viewpoints and the presence of occlusions. We compare the performance in terms of mean average precision (mAP) which is the principal quantitative measure in VOC object detection task [25]. The results demonstrate that our method can boost the detection performance based on the baseline method.

Algorithm 1 Mean field in fully connected CRFs

1: Initialize Q

$$Q_i(x_i) = \frac{1}{Z_i} exp\{-\phi_u(x_i) - \omega_g \phi_g(x_i)\}$$

2: **while** not converged **do**

3: Message passing.

$$\tilde{Q_i}(l) = \sum_{j \neq i} Q_j(l)$$

4: Compatibility transform

$$\hat{Q_i}(x_i) = \sum_{l \in L} \tilde{Q_i}(l) \omega_p \phi_p(l, x_i, r)$$

5: Local update

$$Q_i(x_i) = \frac{1}{Z_i} exp\{-\phi_u(x_i) - \omega_g \phi_g(x_i) - \hat{Q_i}(x_i)\}$$

6: **end while**

4.1 Experiment and Evaluation Details

All of our experiments use the Faster R-CNN method [9] as our baseline. In order to show that our method is insensitive to the stage of object proposal, we utilize two CNN architecture including ZFNet [39] and VGGNet [40] to train the Faster R-CNN system. In addition, we use different traing data to obtain different CNN model. In what follows, we use ZF to denote Faster R-CNN with ZFNet, VGG to denote Faster R-CNN with VGGNet. Pp means the pairwise potentials are added, while Gp means the incorporation of the global potentials. The weight parameters for the pairwise potentials and the global potentials are selected via cross validation. We show the performance in terms of AP for each class on VOC 2007 test.

4.2 Analysis

Firstly, we examine the influence of parameters in our model. Then we show the performance in terms of average precision on the VOC 2007 test.

Parameter Selection In our CRF framework, the parameter ω_p and ω_g in (4) adjust the tradeoff between the local apperance and the contextual information. Here we focus on ω_p which represents the importance of object-level contextual information and other hyperparameters are fixed. The results of average precision for classes bottle, cow and pottedplants are shown in Figure 4. As shown in Figure4, when the parameter ω_p becoms larger, our method achieves better AP at first and then decrease. It demonstrates that there exists a balance for local appearance and contextual coherent constraints among object candidates in the same image.

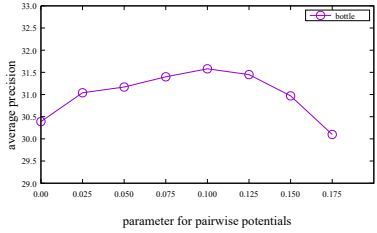

(a) AP for Bottle

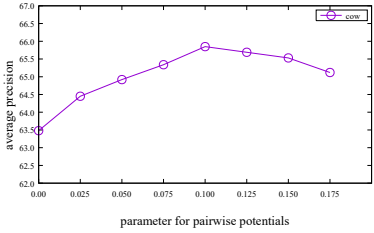

(b) AP for Cow

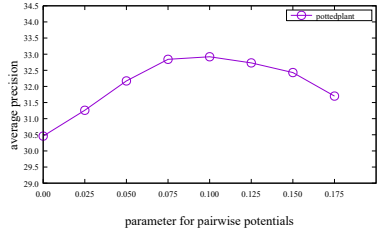

(c) AP for Pottedplant

Fig. 4. AP for three classes with different ω_p

performance on the ZFnet In Table 1, we take ZFnet which is trained on the VOC 2007 trainval dataset as the baseline and our approach obtains 0.87% improvement in terms of mAP. In this experiment, our method yields the best performance over 18 classes. Among these categories, the class bottle is enhanced by 3.4% and class tvmonitor achieves 2.2% improvement. Both pairwise potentials and global potentials support the object detection. Morerover, we can see that the pairwise potentials part plays a more important role because it achieves improvement in 8 calsses compared to 3 classes' enhancement caused by the global potentials. We show the results of ZFnet trained on the union of VOC 2007 and 2012 trainval datasets in Table 2.

Table 1. The result of VOC 2007 test dataset based on ZFnet trained with VOC2007 trainval dataset

class	aero	bike	bird	boat	bottle	bus	car	cat	chair	cow
ZF	64.0	**69.9**	56.6	44.9	30.3	66.8	73.4	71.0	35.6	63.4
ZF+Pp	65.4	69.6	**57.8**	45.6	31.5	67.1	73.6	**71.6**	36.2	**65.8**
ZF+Gp	64.9	69.3	57.0	45.4	33.4	**67.8**	73.7	70.8	36.5	64.5
ZF+Pp+Gp	**65.9**	69.8	57.2	**45.6**	**33.7**	67.2	**73.8**	71.1	**37.3**	65.0
class	table	dog	horse	m-bike	person	plant	sheep	sofa	train	tv
ZF	60.2	65.5	76.7	**70.9**	64.4	30.4	58.0	53.5	72.3	56.6
ZF+Pp	58.3	**66.5**	76.6	70.6	**64.6**	**32.9**	**61.0**	52.6	**72.8**	57.3
ZF+Gp	**61.5**	65.1	76.7	70.8	64.0	31.0	56.6	**54.6**	72.6	58.7
ZF+Pp+Gp	61.4	65.6	**77.2**	70.3	64.1	31.8	59.4	54.2	72.6	**58.8**

Table 2. The result of VOC 2007 test dataset based on ZFnet trained with VOC2007 and VOC2012 trainval dataset

class	aero	bike	bird	boat	bottle	bus	car	cat	chair	cow
ZF	67.8	71.2	59.1	49.8	33.8	71.7	**75.2**	**79.9**	38.4	70.4
ZF+Pp	68.1	70.8	**60.3**	49.5	35.6	**72.1**	74.9	79.3	39.3	**72.3**
ZF+Gp	68.8	70.9	60.0	**50.2**	37.6	71.7	75.1	78.9	39.1	70.4
ZF+Pp+Gp	**69.2**	**71.2**	60.2	48.9	**38.3**	71.7	75.0	79.5	**39.7**	71.8
class	table	dog	horse	m-bike	person	plant	sheep	sofa	train	tv
ZF	58.9	74.1	79.8	72.5	65.0	30.3	66.4	60.6	**72.4**	58.5
ZF+Pp	59.8	74.0	**80.4**	71.7	**65.5**	**33.1**	**67.6**	59.1	72.3	59.2
ZF+Gp	60.3	73.5	79.8	72.6	65.2	32.1	66.5	**61.1**	72.1	59.6
ZF+Pp+Gp	**60.3**	**74.2**	79.5	**73.2**	65.3	32.3	67.5	61.0	72.3	**60.0**

performance on the VGGnet From Table 3, we can see that the result of our method based on VGGnet is 70.18% in terms of mAP which acchives 0.52% improvement. And it outperforms other alternatives on 18 classes out of a total of 20 of them. Furthermore, the result demonstrates that our method have the potential to work similarly well on different object detection methods since it only needs the proposals and the corresponding scores. However, the improvement becomes lower than it achieved by the model based on ZFnet and the reason maybe that the VGGnet CNN model which focuses on candidates themselves are powerful enough to overlook the help from contextual information in some extent. We show the results of VGGnet trained on the union of VOC 2007 and 2012 trainval datasets in Table 4.

Table 3. The result of VOC 2007 test dataset based on VGGnet

class	aero	bike	bird	boat	bottle	bus	car	cat	chair	cow
VGG	69.1	78.8	67.7	**54.8**	49.4	78.1	79.9	**84.6**	50.6	74.2
VGG+Pp	**69.9**	78.8	68.6	54.3	49.3	**78.8**	79.9	84.3	51.2	**77.2**
VGG+Gp	69.1	78.7	67.8	53.7	**51.6**	78.2	**80.0**	84.4	**51.3**	74.8
VGG+Pp+Gp	69.1	**78.8**	**69.6**	53.2	51.0	78.4	79.9	84.5	51.0	76.8
class	table	dog	horse	m-bike	person	plant	sheep	sofa	train	tv
VGG	65.5	81.1	83.6	77.0	75.7	38.4	70.1	66.9	80.6	66.0
VGG+Pp	64.9	81.2	**84.5**	**77.6**	**75.8**	40.5	**71.2**	65.8	80.6	66.3
VGG+Gp	**67.0**	81.2	84.1	76.7	75.6	40.2	70.0	**67.4**	80.7	**67.7**
VGG+Pp+Gp	66.7	**81.2**	84.3	77.5	75.6	**40.6**	70.9	65.7	**80.9**	67.0

4.3 Visualization of More Results

Figures 5 visualizes the performance of our method against Faster R-CNN on some VOC2007 images. In most situations, our method can improve the detection performance. However, the results of the image in the third line which is

Table 4. The result of VOC 2007 test dataset based on VGGnet trained with VOC2007 and VOC2012 trainval dataset

class	aero	bike	bird	boat	bottle	bus	car	cat	chair	cow
VGG	75.4	79.8	74.5	59.9	52.7	82.9	84.6	**88.3**	52.5	79.2
VGG+Pp	75.9	80.0	**75.3**	**59.9**	53.9	**83.0**	84.7	87.9	53.3	**83.4**
VGG+Gp	76.3	80.0	74.6	58.0	55.6	82.3	84.6	87.8	51.3	74.8
VGG+Pp+Gp	**76.3**	**80.3**	75.2	57.8	**56.2**	82.1	**84.9**	87.6	**53.9**	82.2
class	table	dog	horse	m-bike	person	plant	sheep	sofa	train	tv
VGG	66.0	84.8	85.0	**76.8**	76.6	36.8	75.6	**73.1**	81.7	71.5
VGG+Pp	65.7	**85.4**	**85.5**	76.6	**77.0**	40.2	75.9	71.5	**82.2**	**71.5**
VGG+Gp	66.9	84.9	84.0	76.5	76.5	40.3	75.6	72.7	82.1	71.1
VGG+Pp+Gp	**67.1**	85.3	85.2	76.7	76.6	**40.4**	**76.0**	72.4	82.0	71.0

full of tiny boats are worse than the results of ZFnet. Actually, we find that it's because our global contextual part could not recognize the scene as a lake well. And it reminds us there still has very large development space.

5 Conclusions

In this paper, we have proposed an ensamble object detection system which combines the local apperance, the contextual relationships among different objects and the scene context information. Here, we leverage the powerful deep convolutional neural networks to obtain unary potentials for object proposals and extract features representing scene context. In addition, the pairwise potentials which take both semantic and spatial relevance into account for different object proposals are utilized to produce a semantically coherent interpretation of the input image. We fomulate the whole problem in the form of a fully-connected CRF model which can be efficiently solved by a fast mean field inference method. Furthermore, our experimental evaluation has demonstrated that our approach could effectively leverage the contextual information to improve detection accuracy, thus outperforming existing detection techniques on benchmark datasets. In the future, we will devise a better pairwise model based on CNN and incorporate it into an end-to-end framework.

References

1. Carreira, J., Sminchisescu, C.: Cpmc: Automatic object segmentation using constrained parametric min-cuts. Pattern Analysis and Machine Intelligence, IEEE Transactions on **34**(7) (2012) 1312–1328
2. Uijlings, J.R., van de Sande, K.E., Gevers, T., Smeulders, A.W.: Selective search for object recognition. International journal of computer vision **104**(2) (2013) 154–171
3. Arbeláez, P., Pont-Tuset, J., Barron, J., Marques, F., Malik, J.: Multiscale combinatorial grouping. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. (2014) 328–335

Fig. 5. Visualization of hits and misses on VOC 2007 images.In each image, the first and third columns show the results of ZFnet and VGGnet. The results of our contextual model based on ZFnet and VGGnet are displayed in second and forth columns.

4. Krähenbühl, P., Koltun, V.: Geodesic object proposals. In: Computer Vision–ECCV 2014. Springer (2014) 725–739
5. Zitnick, C.L., Dollár, P.: Edge boxes: Locating object proposals from edges. In: Computer Vision–ECCV 2014. Springer (2014) 391–405
6. Felzenszwalb, P.F., Girshick, R.B., McAllester, D., Ramanan, D.: Object detection with discriminatively trained part-based models. Pattern Analysis and Machine Intelligence, IEEE Transactions on **32**(9) (2010) 1627–1645
7. Viola, P., Jones, M.J.: Robust real-time face detection. International journal of computer vision **57**(2) (2004) 137–154
8. Girshick, R., Donahue, J., Darrell, T., Malik, J.: Rich feature hierarchies for accurate object detection and semantic segmentation. In: Computer Vision and Pattern Recognition (CVPR), 2014 IEEE Conference on, IEEE (2014) 580–587
9. Ren, S., He, K., Girshick, R., Sun, J.: Faster r-cnn: Towards real-time object detection with region proposal networks. In: Advances in Neural Information Processing Systems. (2015) 91–99
10. Biederman, I.: Perceiving real-world scenes. Science **177**(4043) (1972) 77–80
11. Biederman, I., Mezzanotte, R.J., Rabinowitz, J.C.: Scene perception: Detecting and judging objects undergoing relational violations. Cognitive psychology **14**(2)

(1982) 143–177
12. Torralba, A.: Contextual priming for object detection. International journal of computer vision **53**(2) (2003) 169–191
13. Crandall, D.J., Huttenlocher, D.P.: Composite models of objects and scenes for category recognition. In: Computer Vision and Pattern Recognition, 2007. CVPR'07. IEEE Conference on, IEEE (2007) 1–8
14. Rabinovich, A., Vedaldi, A., Galleguillos, C., Wiewiora, E., Belongie, S.: Objects in context. In: Computer vision, 2007. ICCV 2007. IEEE 11th international conference on, IEEE (2007) 1–8
15. Tu, Z.: Auto-context and its application to high-level vision tasks. In: Computer Vision and Pattern Recognition, 2008. CVPR 2008. IEEE Conference on, IEEE (2008) 1–8
16. Galleguillos, C., Rabinovich, A., Belongie, S.: Object categorization using co-occurrence, location and appearance. In: Computer Vision and Pattern Recognition, 2008. CVPR 2008. IEEE Conference on, IEEE (2008) 1–8
17. Song, Z., Chen, Q., Huang, Z., Hua, Y., Yan, S.: Contextualizing object detection and classification. In: Computer Vision and Pattern Recognition (CVPR), 2011 IEEE Conference on, IEEE (2011) 1585–1592
18. Choi, M.J., Lim, J.J., Torralba, A., Willsky, A.S.: Exploiting hierarchical context on a large database of object categories. In: Computer vision and pattern recognition (CVPR), 2010 IEEE conference on, IEEE (2010) 129–136
19. Desai, C., Ramanan, D., Fowlkes, C.C.: Discriminative models for multi-class object layout. International journal of computer vision **95**(1) (2011) 1–12
20. Mottaghi, R., Chen, X., Liu, X., Cho, N.G., Lee, S.W., Fidler, S., Urtasun, R., Yuille, A.: The role of context for object detection and semantic segmentation in the wild. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. (2014) 891–898
21. Dalal, N., Triggs, B.: Histograms of oriented gradients for human detection. In: Computer Vision and Pattern Recognition, 2005. CVPR 2005. IEEE Computer Society Conference on. Volume 1., IEEE (2005) 886–893
22. Krizhevsky, A., Sutskever, I., Hinton, G.E.: Imagenet classification with deep convolutional neural networks. In: Advances in neural information processing systems. (2012) 1097–1105
23. Krähenbühl, P., Koltun, V.: Efficient inference in fully connected crfs with gaussian edge potentials. arXiv preprint arXiv:1210.5644 (2012)
24. Hayder, Z., Salzmann, M., He, X.: Object co-detection via efficient inference in a fully-connected crf. In: Computer Vision–ECCV 2014. Springer (2014) 330–345
25. Everingham, M., Van Gool, L., Williams, C., Winn, J., Zisserman, A.: The pascal visual object classes challenge 2007 (voc 2007) results (2007) (2008)
26. Szegedy, C., Toshev, A., Erhan, D.: Deep neural networks for object detection. In: Advances in Neural Information Processing Systems. (2013) 2553–2561
27. Sermanet, P., Eigen, D., Zhang, X., Mathieu, M., Fergus, R., LeCun, Y.: Overfeat: Integrated recognition, localization and detection using convolutional networks. arXiv preprint arXiv:1312.6229 (2013)
28. Erhan, D., Szegedy, C., Toshev, A., Anguelov, D.: Scalable object detection using deep neural networks. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. (2014) 2147–2154
29. He, K., Zhang, X., Ren, S., Sun, J.: Spatial pyramid pooling in deep convolutional networks for visual recognition. Pattern Analysis and Machine Intelligence, IEEE Transactions on **37**(9) (2015) 1904–1916

30. Girshick, R.: Fast r-cnn. In: Proceedings of the IEEE International Conference on Computer Vision. (2015) 1440–1448
31. Bell, S., Zitnick, C.L., Bala, K., Girshick, R.: Inside-outside net: Detecting objects in context with skip pooling and recurrent neural networks. arXiv preprint arXiv:1512.04143 (2015)
32. Liu, S., Lu, C., Jia, J.: Box aggregation for proposal decimation: Last mile of object detection. In: Proceedings of the IEEE International Conference on Computer Vision. (2015) 2569–2577
33. Gidaris, S., Komodakis, N.: Locnet: Improving localization accuracy for object detection. arXiv preprint arXiv:1511.07763 (2015)
34. Redmon, J., Divvala, S., Girshick, R., Farhadi, A.: You only look once: Unified, real-time object detection. arXiv preprint arXiv:1506.02640 (2015)
35. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. arXiv preprint arXiv:1512.03385 (2015)
36. Yao, B., Fei-Fei, L.: Modeling mutual context of object and human pose in human-object interaction activities. In: Computer Vision and Pattern Recognition (CVPR), 2010 IEEE Conference on, IEEE (2010) 17–24
37. Li, B., Wu, T., Zhu, S.C.: Integrating context and occlusion for car detection by hierarchical and-or model. In: Computer Vision–ECCV 2014. Springer (2014) 652–667
38. Vu, T.H., Osokin, A., Laptev, I.: Context-aware cnns for person head detection. In: Proceedings of the IEEE International Conference on Computer Vision. (2015) 2893–2901
39. Zeiler, M.D., Fergus, R.: Visualizing and understanding convolutional networks. In: Computer vision–ECCV 2014. Springer (2014) 818–833
40. Simonyan, K., Zisserman, A.: Very deep convolutional networks for large-scale image recognition. arXiv preprint arXiv:1409.1556 (2014)
41. LeCun, Y., Boser, B., Denker, J.S., Henderson, D., Howard, R.E., Hubbard, W., Jackel, L.D.: Backpropagation applied to handwritten zip code recognition. Neural computation **1**(4) (1989) 541–551
42. Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L.: Imagenet: A large-scale hierarchical image database. In: Computer Vision and Pattern Recognition, 2009. CVPR 2009. IEEE Conference on, IEEE (2009) 248–255
43. Zhou, B., Khosla, A., Lapedriza, A., Torralba, A., Oliva, A.: Places2: A large-scale database for scene understanding. Arxiv, 2015 (coming soon) (2015)
44. Zhou, B., Lapedriza, A., Xiao, J., Torralba, A., Oliva, A.: Learning deep features for scene recognition using places database. In: Advances in neural information processing systems. (2014) 487–495